Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L3	293	(715/511).CCLS.	US-PGPUB; USPAT	OR	OFF	2005/03/18 13:50
L4	871	(715/530).CCLS.	US-PGPUB; USPAT	OR	OFF	2005/03/18 13:49
L5	141	(715/509).CCLS.	US-PGPUB; USPAT	OR	OFF	2005/03/18 13:49
L6	1166	stor$4 with version with table	US-PGPUB; USPAT	OR	ON	2005/03/18 13:50
L7	55	stor$4 with (new$4 adj version) with table	US-PGPUB; USPAT	OR	ON	2005/03/18 13:50
L8	40	stor$4 with (old$4 adj version) with table	US-PGPUB; USPAT	OR	ON	2005/03/18 13:51
L9	9	7 and 8	US-PGPUB; USPAT	OR	ON	2005/03/18 13:51
L18	16195	version with (table or database)	US-PGPUB; USPAT; EPO	OR	ON	2005/03/18 13:52
L19	3430	L18 and (new$4 with version)	US-PGPUB; USPAT; EPO	OR	ON	2005/03/18 13:52
L20	107	L19 and (new$4 with version with first with (table or database))	US-PGPUB; USPAT; EPO	OR	ON	2005/03/18 13:52
L21	426	old$4 with version with (table or database)	US-PGPUB; USPAT; EPO	OR	ON	2005/03/18 13:52
L22	23	L20 and L21	US-PGPUB; USPAT; EPO	OR	ON	2005/03/18 13:53
L23	408	sql with column with query	US-PGPUB; USPAT	OR	ON	2005/03/18 13:53
L24	7309	obtain$4 adj2 column	US-PGPUB; USPAT	OR	ON	2005/03/18 13:53
L25	17	L23 and L24	US-PGPUB; USPAT	OR	ON	2005/03/18 13:54
L26	52	determin$4 with version with number with location	US-PGPUB; USPAT	OR	ON	2005/03/18 13:54

Terms used **store data version tables** Found **85,637** of **151,219**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 Query and view processing: On querying versions of multiversion data warehouse

Tadeusz Morzy, Robert Wrembel

November 2004 **Proceedings of the 7th ACM international workshop on Data warehousing and OLAP**

Full text available: pdf(244.86 KB) Additional Information: full citation, abstract, references, index terms

A data warehouse (DW) is fed with data that come from external data sources that are production systems. External data sources, which are usually autonomous, often change not only their content but also their structure. The evolution of external data sources has to be reflected in a DW, that uses the sources. Traditional DW systems offer a limited support for handling dynamics in their structure and content. A promising approach to handling changes in DW structure and content is based on a mu ...

Keywords: data versioning, data warehouse, metadata, multiversion query, schema versioning

2 Database theory, technology and applications (DTTA): Creation and management of versions in multiversion data warehouse

Bartosz B□bel, Johann Eder, Christian Koncilia, Tadeusz Morzy, Robert Wrembel

March 2004 **Proceedings of the 2004 ACM symposium on Applied computing**

Full text available: pdf(516.99 KB) Additional Information: full citation, abstract, references, index terms

A data warehouse (DW) provides an information for analytical processing, decision making, and data mining tools. On the one hand, the structure and content of a data warehouse reflects a real world, i.e. data stored in a DW come from real production systems. On the other hand, a DW and its tools may be used for predicting trends and simulating a virtual business scenarios. This activity is often called the what-if analysis. Traditional DW systems have static structure of their schemas and relati ...

Keywords: data warehouse, integrity constraints, versioning

3 Industrial sessions: beyond relational tables: Coordinating backup/recovery and data consistency between database and file systems

Suparna Bhattacharya, C. Mohan, Karen W. Brannon, Inderpal Narang, Hui-I Hsiao, Mahadevan Subramanian

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on**

Management of data

Full text available: pdf(1.44 MB) Additional Information: full citation, abstract, references, index terms

Managing a combined store consisting of database data and file data in a robust and consistent manner is a challenge for database systems and content management systems. In such a hybrid system, images, videos, engineering drawings, etc. are stored as files on a file server while meta-data referencing/indexing such files is created and stored in a relational database to take advantage of efficient search. In this paper we describe solutions for two potentially problematic aspects of such a data ...

Keywords: DB2, content management, database backup, database recovery, datalinks

4 Using Applications of Data Versioning in Database Application Development

Ramkrishna Chatterjee, Gopalan Arun, Sanjay Agarwal, Ben Speckhard, Ramesh Vasudevan

May 2004 **Proceedings of the 26th International Conference on Software Engineering - Volume 00**

Full text available: pdf(166.57 KB) Additional Information: full citation, abstract

Database applications such as enterprise resource planning systems and customer relationship management systems are widely used software systems. Development and testing of database applications is difficult because the program execution depends on the persistent state stored in the database. In this paper we show that how versioning of the persistent data stored in the database can solve some critical problems in the development and testing of database applications can be solved by vers ...

5 A linear-time scheme for version reconstruction

Lin Yu, Daniel J. Rosenkrantz

May 1994 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 16 Issue 3

Full text available: pdf(1.47 MB) Additional Information: full citation, abstract, references, citings, index terms, review

An efficient scheme to store and reconstruct versions of sequential files is presented. The reconstruction scheme involves building a data structure representing a complete version, and then successively modifying this data structure by applying a sequence of specially formatted differential files to it. Each application of a differential file produces a representation of an intermediate version, with the final data structure representing the requested version. The scheme uses a ...

Keywords: data structures, database systems, differential files, document preparation, software systems, textual objects, version control

6 Effective fine-grain synchronization for automatically parallelized programs using optimistic synchronization primitives

Martin C. Rinard

November 1999 **ACM Transactions on Computer Systems (TOCS)**, Volume 17 Issue 4

Full text available: pdf(637.69 KB) Additional Information: full citation, abstract, references, citings, index terms, review

This article presents our experience using optimistic synchronization to implement fine-grain atomic operations in the context of a parallelizing compiler for irregular, object-based computations. Our experience shows that the synchronization requirements of these programs differ significantly from those of traditional parallel computations, which use loop nests to access dense matrices using affine access functions. In addition to coarse-grain barrier synchronization, our irregular comput ...

Keywords: atomic operations commutativity analysis, optimistic synchronization, parallel computing, parallelizing compilers, synchronization

7 The BUCKY object-relational benchmark

Michael J. Carey, David J. DeWitt, Jeffrey F. Naughton, Mohammad Asgarian, Paul Brown, Johannes E. Gehrke, Dhaval N. Shah

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2

Full text available: pdf(1.48 MB) Additional Information: full citation, abstract, references, citings, index terms

According to various trade journals and corporate marketing machines, we are now on the verge of a revolution—the object-relational database revolution. Since we believe that no one should face a revolution without appropriate armaments, this paper presents BUCKY, a new benchmark for object-relational database systems. BUCKY is a query-oriented benchmark that tests many of the key features offered by object-relational systems, including row types and inheritance, references and path e ...

8 SmartFiles: an OO approach to data file interoperability

Matthew Haines, Piyush Mehrotra, John Van Rosendale

October 1995 **ACM SIGPLAN Notices , Proceedings of the tenth annual conference on Object-oriented programming systems, languages, and applications**, Volume 30 Issue 10

Full text available: pdf(1.59 MB) Additional Information: full citation, abstract, references, citings, index terms

Data files for scientific and engineering codes typically consist of a series of raw data values whose description is buried in the programs that interact with these files. In this situation, making even minor changes in the file structure or sharing files between programs (interoperability) can only be done after careful examination of the data files and the I/O statements of the programs interacting with this file. In short, scientific data files lack self-description, and other self-describin ...

9 P1: "Yes, but does it scale?": practical considerations for database-driven information systems

John Russell

October 2001 **Proceedings of the 19th annual international conference on Computer documentation**

Full text available: pdf(231.31 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper explores the process of designing and implementing a database-driven system of online documentation, and putting it live on the web for customers to use. Using real-life examples, it discusses practical considerations for balancing performance, scalability, and reliability.

Keywords: Oracle, automation, categorization, database, performance, reliability, scalability, web services

10 Profiling a parallel language based on fine-grained communication

Bjoern Haake, Klaus E. Schauser, Chris Scheiman

November 1996 **Proceedings of the 1996 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available: pdf(198.94 KB) Additional Information: full citation, abstract, references, index terms

Fine tuning the performance of large parallel programs is a very difficult task. A profiling tool can provide detailed insight into the utilization and communication of the different processors, which helps identify performance bottlenecks. In this paper we present a profiler for the fine-grained parallel programming language Split-C, which provides a simple global address space memory model. As our experience shows, it is much more challenging to profile programs that make use of efficient ...

11 An experimental analysis of the performance of fourth generation tools on PCs

Victor M. Matos, Paul J. Jalics

November 1989 **Communications of the ACM**, Volume 32 Issue 11

Full text available: pdf(1.52 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The performance of several Fourth Generation Language (4GL) tools is analyzed empirically and compared with equivalent programs written in the third generation COBOL programming language. A set of performance benchmarks consisting of thirteen separate functions is presented which encompasses the areas of simulating the operators of the relational algebra, accessing records in the database, and updating the database. This serves as a baseline for comparing the various 4GL systems.

12 Satellite-based information services: Cost based data dissemination in satellite networks

Bo Xu, Ouri Wolfson, Sam Chamberlain, Naphtali Rishe

January 2002 **Mobile Networks and Applications**, Volume 7 Issue 1

Full text available: pdf(312.11 KB) Additional Information: full citation, abstract, references, citings, index terms

We consider the problem of data dissemination in a broadcast network. In contrast to previously studied models, broadcasting is among peers, rather than client server. Such a model represents, for example, satellite communication among widely distributed nodes, sensor networks, and mobile ad hoc networks. We introduce a cost model for data dissemination in peer to peer broadcast networks. The model quantifies the tradeoff between the inconsistency of the data, and its transmission cost; the tran ...

Keywords: data replication, distributed databases, satellite networks

13 Are bitvectors optimal?

H. Buhrman, P. B. Miltersen, J. Radhakrishnan, S. Venkatesh

May 1999 **Proceedings of the thirty-second annual ACM symposium on Theory of computing**

Full text available: pdf(1.01 MB) Additional Information: full citation, references, citings, index terms

14 Database performance in the real world: TPC-D and SAP R/3

Joachen Doppelhammer, Thomas Höppler, Alfons Kemper, Donald Kossmann

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2

Full text available: pdf(1.54 MB) Additional Information: full citation, abstract, references, citings, index terms

Traditionally, database systems have been evaluated in isolation on the basis of standardized benchmarks (e.g., Wisconsin, TPC-C, TPC-D). We argue that very often such a performance analysis does not reflect the actual use of the DBMSs in the "real world." End users typically don't access a stand-alone database system; rather they use a comprehensive application system, in which the database system constitutes an integrated

component. In order to derive performance evalu ...

15 Experiments with an ocean circulation model on CEDAR
L. DeRose, K. Gallivan, E. Gallopoulos
August 1992 **Proceedings of the 6th international conference on Supercomputing**
Full text available: pdf(1.41 MB) Additional Information: full citation, abstract, references, index terms

We present the design of the GFDL ocean circulation model as adapted for simulations of the Mediterranean basis for the Cedar multicluster architecture. The model simulates the basic aspects of large-scale, baroclinic ocean circulation, including treatment of irregular bottom topography. The data and computational mapping strategies and their effect on the design are discussed. The code was parametrized to offer several choices for data partitionings of the computational domain, for placeme ...

16 The logical disk: a new approach to improving file systems
Wiebren de Jonge, M. Frans Kaashoek, Wilson C. Hsieh
December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles**, Volume 27 Issue 5
Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, citings, index terms

The Logical Disk (LD) defines a new interface to disk storage that separates file management and disk management by using logical block numbers and block lists. The LD interface is designed to support multiple file systems and to allow multiple implementations, both of which are important given the increasing use of kernels that support multiple operating system personalities.A log-structured implementation of LD (LLD) demonstrates that LD can be implemented efficiently. LLD adds about 5% to 10% ...

Keywords: MINIX, UNIX, disk storage management, file system organization, file system performance, high write performance, log-structured file system, logical disk

17 Physical design: Handling big dimensions in distributed data warehouses using the DWS technique
Marco Costa, Henrique Madeira
November 2004 **Proceedings of the 7th ACM international workshop on Data warehousing and OLAP**
Full text available: pdf(288.33 KB) Additional Information: full citation, abstract, references, index terms

The DWS (Data Warehouse Striping) technique allows the distribution of large data warehouses through a cluster of computers. The data partitioning approach partition the facts tables through all nodes and replicates the dimension tables. The replication of the dimension tables creates a limitation to the applicability of the DWS technique to data warehouses with big dimensions. This paper proposes a strategy to handle large dimensions in a distributed DWS system and evaluates the proposed str ...

Keywords: data warehousing, distributed query execution

18 Compiler-directed run-time monitoring of program data access
Chen Ding, Yutao Zhong
June 2002 **ACM SIGPLAN Notices , Proceedings of the workshop on Memory system performance**, Volume 38 Issue 2 supplement
Full text available: pdf(1.40 MB) Additional Information: full citation, abstract, references, citings

Accurate run-time analysis has been expensive for complex programs, in part because most

methods perform on all a data. Some applications require only partial reorganization. An example of this is off-loading infrequently used data from a mobile device. Complete monitoring is not necessary because not all accesses can reach the displaced data. To support partial monitoring, this paper presents a framework that includes a source-to-source C compiler and a run-time monitor. The compiler inserts ru ...

19 Making data structures confluently persistent

Amos Fiat, Haim Kaplan

January 2001 **Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms**

Full text available: pdf(924.65 KB) Additional Information: full citation, abstract, references, citings, index terms

We address a longstanding open problem of [8, 7], and present a general transformation that takes any data structure and transforms it to a confluently persistent data structure. We model this general problem using the concepts of a version DAG (Directed Acyclic Graph) and an instantiation of a version DAG. We introduce the concept of the effective depth of a vertex in the version DAG and use it to derive information theoretic lower bounds on the space expansion of any such transformation for ...

20 Continuous queries over append-only databases

Douglas Terry, David Goldberg, David Nichols, Brian Oki

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data**, Volume 21 Issue 2

Full text available: pdf(1.19 MB) Additional Information: full citation, abstract, references, citings, index terms

In a database to which data is continually added, users may wish to issue a permanent query and be notified whenever data matches the query. If such continuous queries examine only single records, this can be implemented by examining each record as it arrives. This is very efficient because only the incoming record needs to be scanned. This simple approach does not work for queries involving joins or time. The Tapestry system allows users to issue such queries over a databa ...

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

	Document ID	Issue Date	Title	Current OR
1	US 20040177058 A1	20040909	Navigation of the content space of a document set	707/1
2	US 20040153469 A1	20040805	Database comparator	707/101
3	US 20040098715 A1	20040520	Over the air mobile device software management	717/173
4	US 20030105737 A1	20030605	Method for the online transformation of bulk data	707/1
5	US 20030061245 A1	20030327	Implementing versioning support for data using a two-table approach that maximizes database efficiency	707/203
6	US 20020143764 A1	20021003	Data management system and method for intercepting and changing database instructions between a database back end and an application front end	707/8
7	US 20020103815 A1	20020801	High speed data updates implemented in an information storage and retrieval system	707/203
8	US 6636802 B1	20031021	Data structure of digital map file	701/208
9	US 6560776 B1	20030506	Software installation verification tool	717/176

	Document ID	Issue Date	Title	Current OR
10	US 6560601 B1	20030506	Database transaction with locking in two phases and multiple versions of objects	707/8
11	US 6519601 B1	20030211	Relational database compiled/stored on a memory structure providing improved access through use of redundant representation of data	707/100
12	US 6263506 B1	20010717	Data transmission and reception device and system, data transmission method and parameter setting method for data reception device	725/116
13	US 6223344 B1	20010424	Apparatus and method for versioning persistent objects	717/170
14	US 6216136 B1	20010410	Method for performing complicated schema changes within a database	707/203
15	US 6085036 A	20000704	Data base structure and management	717/120
16	US 6044205 A	20000328	Communications system for transferring information between memories according to processes transferred with the information	709/201

	Document ID	Issue Date	Title	Current OR
17	US 6026412 A	20000215	Interaction between application of a log and maintenance of a table that maps record identifiers during online reorganization of a database	707/200
18	US 5758156 A	19980526	Method and apparatus of testing programs	713/100
19	US 5721915 A	19980224	Interaction between application of a log and maintenance of a table that maps record identifiers during online reorganization of a database	707/200
20	US 5671428 A	19970923	Collaborative document processing system with version and comment management	715/751
21	US 5404525 A	19950404	Efficient method router that supports multiple simultaneous object versions	717/108
22	US 5212789 A	19930518	Method and apparatus for updating application databases used in a distributed transaction processing environment	707/8

	Document ID	Issue Date	Title	Current OR
23	US 5089954 A	19920218	Method for handling conversational transactions in a distributed processing environment	707/10